

Fringe fitting for coherent integrations with the NPOI.

A. M. Jorgensen,^a D. Mozurkewich,^b T. Armstrong,^c R. Hindsley,^c T. Pauls,^c C. Gilbreath,^c S. Restaino^d

^aLos Alamos National Laboratory, Los Alamos, NM, USA;

^bSeabrook Engineering, Seabrook, MD, USA;

^cNaval Research Laboratory, Washington, DC, USA;

ABSTRACT

We have developed a method for performing long coherent integrations with the Navy Prototype Optical Interferometer (NPOI), which is based on fitting a model fringe pattern to the NPOI data frames. The procedure is quite computationally intensive, but gives a better estimation of the phase than the conventional method of location the peak of the group delay power. We mention briefly some of the most important past work on coherent integration, and then describe our method. We conclude that the fitting approach produces a phase with fewer outliers than the Fourier-transform group delay approach. We show how the instrumental squared visibility varies as a function of the fringe model used, and show that it provides a better SNR than the FT method. The phase determination will always be imperfect, and thus cause a reduction in the visibility amplitude relative to the true instrumental visibility. We illustrate a method for calibrating the visibility amplitude. With long coherent integrations the phase of the visibility can be extracted. We show examples of visibility phases and how to correct them for phase variations in the instrument. Finally, we illustrate a method for measuring stellar diameters very precisely, to one part in at least several hundred.

Keywords: Optical interferometry, coherent integration, NPOI, baseline phases, diameter measurements

1. INTRODUCTION

In order to advance the field of optical interferometry to the sophistication that has been achieved in radio interferometry it is necessary to apply the four techniques of (1) coherent integrations, (2) baseline bootstrapping, (3) baseline phase information, (4) use of visibility information on baselines with very low visibilities. Once all of these are achieved, a range of new questions can be answered, and true imaging becomes possible. In this paper we demonstrate the use of all 4 of these techniques. Coherent integration and bootstrapping are the key techniques which enable us to obtain single baseline phase information and allows us to measure visibilities on resolved baseline. The idea of using coherent integrations to improve the signal-to-noise ratio is not new. Colavita (1999)¹ (and references therein) discuss fringe tracking using the white-light phase of the target, and the improvement in signal-to-noise ratio. Meisner (1998,2000)^{2,3} remonstrated a probabilistic method for determining the fringe phase. He also points out one of the advantages of coherent integration; that the complex visibility (both amplitude and phase) can be obtained directly; And that this phase can be obtained for a low-visibility baseline by using the phase tracked on higher-visibility baselines, i.e. bootstrapping.

A good method for performing very long coherent integrations is through post-processing a large number of individual short coherent integrations to form a single long integration. Quirrenbach (1994)⁴ demonstrated how to coherently integrate in this way in a narrow spectral channel by using the phase in a broader channel as a reference. It is also possible to obtain the fringe phase without using a reference object on the sky. When multiple spectral channels are present the group-delay method can be employed. It takes advantage of the fact that the

Further author information: (Send correspondence to Anders M. Jorgensen)
 Anders. M. Jorgensen: e-mail: ajorg@lanl.gov, telephone: +1-505-667-0944
 D. M. Mozurkewich: e-mail: dave@mozurkewich.com, telephone: +1-301-459-3375
 T. Armstrong: e-mail: tom.armstrong@nrl.navy.mil, telephone: +1-202-767-0669
 R. Hindsley: e-mail: hindsley@nrl.navy.mil
 T. Pauls: pauls@nrl.navy.mil, telephone: +1-202-767-0171

Report Documentation Page			Form Approved OMB No. 0704-0188	
<p>Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p>				
1. REPORT DATE 2004	2. REPORT TYPE	3. DATES COVERED 00-00-2004 to 00-00-2004		
4. TITLE AND SUBTITLE Fringe fitting for coherent integrations with the NPOI.		5a. CONTRACT NUMBER		
		5b. GRANT NUMBER		
		5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)		5d. PROJECT NUMBER		
		5e. TASK NUMBER		
		5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, 4555 Overlook Avenue SW, Washington, DC, 20375		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)		
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT We have developed a method for performing long coherent integrations with the Navy Prototype Optical Interferometer (NPOI), which is based on fitting a model fringe pattern to the NPOI data frames. The procedure is quite computationally intensive, but gives a better estimation of the phase than the conventional method of location the peak of the group delay power. We mention briefly some of the most important past work on coherent integration, and then describe our method. We conclude that the fitting approach produces a phase with fewer outliers than the Fourier-transform group delay approach. We show how the instrumental squared visibility varies as a function of the fringe model used, and show that it provides a better SNR than the FT method. The phase determination will always be imperfect, and thus cause a reduction in the visibility amplitude relative to the true instrumental visibility. We illustrate a method for calibrating the visibility amplitude. With long coherent integrations the phase of the visibility can be extracted. We show examples of visibility phases and how to correct them for phase variations in the instrument. Finally, we illustrate a method for measuring stellar diameters very precisely, to one part in at least several hundred.				
15. SUBJECT TERMS				
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified	18. NUMBER OF PAGES 12	19a. NAME OF RESPONSIBLE PERSON

fringe phases across spectral channels are related, to first order via just two parameters - the fringe tracking error and the phase at zero fringe tracking error. Hummel et al. (2003)⁵ discuss performing coherent integrations by determining the phase via the group-delay tracking method and integrating on low-visibility baselines using bootstrapping. They demonstrate how this improves the SNR and apply the technique to some limb-darkening measurements.

Lane et al. (2003)⁶ describe a different approach to performing coherent integrations of 100's of ms. They stabilize fringes on a nearby brighter star, thus achieving fairly good stability on the target star. Wittkowski and Hummel (2003)⁷ used coherent integration and baseline bootstrapping with other techniques to measure stellar intensity profiles.

In this paper we demonstrate a method for performing long coherent integrations, by using a fringe fitting method to determine the fringe tracking error and phase, instead of the group-delay method.

2. MOTIVATION

The conventional method for analyzing optical interferometry data uses the approach of computing the fringe power spectrum for individual frames, and then averaging many such power spectra. The averaging of many frames is necessary because the count-rate in a single frame is so low that the uncertainty on the corresponding V^2 is unusably large. The mathematical process for computing V^2 from a single frame is

$$V_K(\lambda)^2 = \frac{4(X_K(\lambda)^2 + Y_K(\lambda)^2 - B(\lambda))}{N^2}, \quad (1)$$

where

$$X_K(\lambda) = \sum_{i=0}^{N_B} C_i(\lambda) \cos\left(\frac{2\pi K i}{N_B}\right) \quad Y_K(\lambda) = \sum_{i=0}^{N_B} C_i(\lambda) \sin\left(\frac{2\pi K i}{N_B}\right) \quad (2)$$

are the cosine and sine transforms of data from a single baseline in a single wavelength band in a single integration. B is the bias which comes about because of the squaring of the noisy quantities X_K and Y_K . For pure Poisson noise, $B = N$. Unfortunately, the noise is often not purely Poisson. There are several fundamental problems with this approach. Firstly, the squaring process (or, really the square modulus process) discards phase information. Secondly, the squaring and subsequent bias subtraction causes a large uncertainty in V^2 in cases where X and Y are small, i.e. situations where the visibility is small or the count rates are small. These also happen to be the most interesting situation, namely resolving baselines on faint stars. For short integrations where the count rate, N , is very small, X_K and Y_K and B are all very small, and a precise determination of B is imperative to getting a good determination of V^2 . If on the other hand we could integrate for longer periods of time so that X_K , Y_K , and B are large, then it becomes less important to determine B accurately.

For these reasons, in order to use data from baselines with low visibilities it is necessary to perform long coherent integrations. To do this it is either necessary to implement a system in hardware which tracks the phase in real time, or to combine many short exposures after the fact with phase corrections. One approach to obtaining the phase in a multi-spectral system such as the NPOI is to use the Fourier-transform group-delay (FT) method [see e.g. Hummel et al. (2003)⁵]. However, the FT method imposes a number of implicit assumptions; (a) That NV is constant with wavelength. (b) That the spectral channels are linearly spaced in wavenumber. (c) That the statistics is either Poisson or Gaussian. (d) Finally, the FT method is only an optimal estimator of the fringe amplitude if there is an integral number of fringe periods in the data set. Since we do not expect all of these assumptions to be fulfilled with the NPOI, or indeed with most optical interferometers, it is at least worth looking into other methods for determining the phase. We intend to look at the method of fitting a model of the fringe pattern to individual frames of data, as well as fitting a time-dependent model of the fringe pattern to a series of frames.

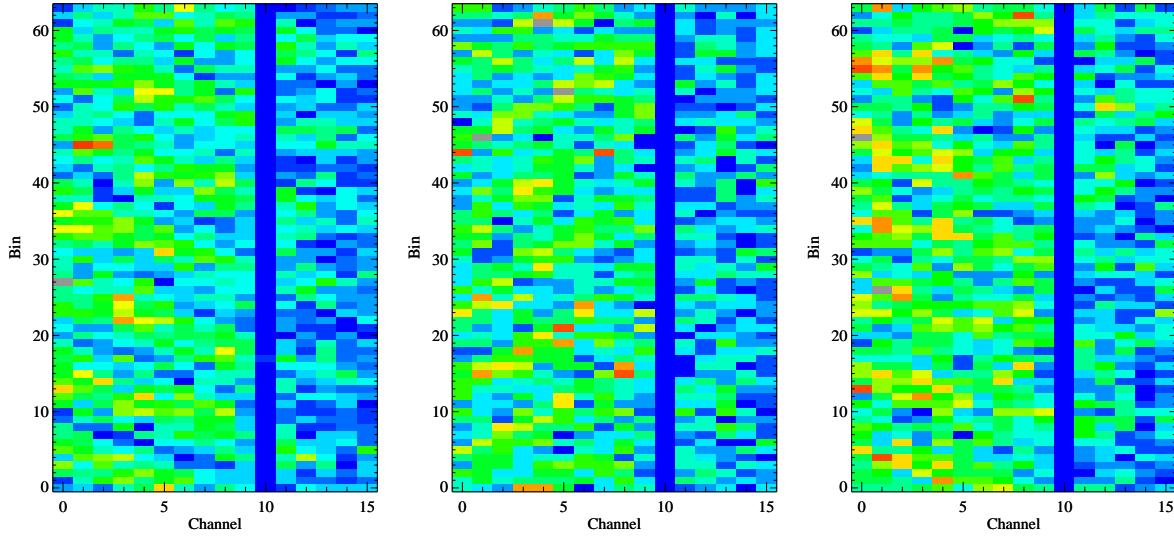


Figure 1. Three of the 15000 data frames that were recorded on the bright star. The frames show intensity - number of counts in 2 ms - as a color code. From faintest to brightest the sequence is dark blue, light blue, green, yellow, orange, red. The horizontal axis enumerates the wavelength channels with channel 0 being 852 nm, and channel 15 being 557 nm. The vertical axis is the stroke.

3. METHOD

We have $X + iY$ for each of a large number of frames, and we wish to determine a single $X + iY$ for all the frames. Before we add them all up we need to first rotate them to a common reference. This can be written as

$$(X + iY)_t = \sum_{i=1}^N (X + iY)_i e^{-i\phi_i} \quad (3)$$

This leaves the problem of determining ϕ_i . If only a single wavelength channel is available, the phase determination is ambiguous. If there are multiple wavelength channels available, their phase offsets are, to first order, related by

$$\phi_i(\lambda) = \phi_{i,0} + 2\pi \frac{d_i}{\lambda}. \quad (4)$$

We are thus down to determining two parameters, the fringe tracking error, d_i , and the fringe phase at zero fringe tracking error, ϕ_i , for each frame. In-fact, we have to do this for each fringe frequency, k , in each frame, so we label these quantities $d_k(i)$, and $\phi_k(i)$.

3.1. NPOI data format

For this paper we used two scans of 30 seconds each, recorded on October 18, 2003. (For a detailed introduction to the NPOI please read Armstrong et al. (1998)⁸). One was a program star, and another was a calibration star. For the brighter program star the average count rate was 4.7 counts in each of 64 by 16 pixels. For the fainter calibration star the count rate was 0.31 in the same number of pixels. The data set contains fringes recorded for three baselines. Fringe frequency 2 is for the longest baseline, frequency 6 is for the shortest baseline, and frequency 8 is for the intermediate baseline. Figure 1 show three 2 ms frames of data for the program star. The wavelengths decrease from left to right from 852 nm to 557 nm. The vertical scale is the stroke, and is adjusted separately for each wavelength channel such that an integer number of fringes is measured at each wavelength. In each of the three frames we can clearly see the $K = 6$ fringes. However, the $K = 2$ and $K = 8$ fringes are too faint to be seen with the eye.

3.2. The single-frame fringe function

We model the phase error for the individual channels as a constant term plus a term that varies inversely with wavelength (see e.g. Quirrenbach et al. (1994)⁴), as

$$\phi = \phi_k + 2\pi \frac{d_k}{\lambda} \quad (5)$$

The fringes in the NPOI frames can be described by the following function (See Born (1980)⁹, Chapter 7 for more details),

$$I(\lambda, b) = I_0(\lambda)[1 + 2 \sum_{k=1}^K V(\lambda, k) \cos(\phi_k + 2\pi(\frac{d_k}{\lambda} + f_k \frac{b - \frac{B-1}{2}}{B}))]. \quad (6)$$

where $I(\lambda, b)$ is the intensity in a channel/bin pixel. There are 16 channels, and 64 bins. $I_0(\lambda)$ is the mean intensity in a channel, K is the number of fringe frequencies, ϕ_k is the fringe phase, d_k is the fringe tracking error, f_k is the fringe frequency, and $B = 64$, the number of bins in a channel.

3.3. Using multiple frames

Equation 6 can be used to calculate the fringe parameters of phase, ϕ_k , and delay, d_k , which in turn can be used to calculate the phase (via Equation 5) needed in Equation 3. Often the count rates in a frame are so low that ϕ_k and d_k cannot easily be determined unambiguously. Firstly, there is noise in the determination, and secondly sometimes secondary minima of the cost function (or maxima of the group-delay power spectrum) may be picked instead of the desired one. But we know that the parameters ϕ_k and d_k are not independent from one frame to the next. We can take advantage of this by modeling ϕ_k and d_k as a function of time and then fit multiple frames simultaneously to improve the accuracy of the determination. For this paper we have chosen to model ϕ_k and d_k as low-order polynomials,

$$d_k(t) = \sum_{i=0}^{N_k} d_{k,i}(t - t_0)^i \quad \phi_k(t) = \sum_{i=0}^{N_k} \phi_{k,i}(t - t_0)^i \quad (7)$$

3.4. The cost function

The cost function is a mathematical expression that defines the goodness of the fit. It is chosen such that minimizing it creates an optimal fit between data and model when the cost function is minimized. Ideally, the cost function should closely match the statistical distribution of the data. The most commonly used cost function is χ^2 , which is the sum of the square of differences between each data pixel and the corresponding model value, divided by the number of data pixels minus the number of parameters. This cost function is best when the statistics are Gaussian. Another common cost function is the total absolute deviation. This is the sum of the absolute values of difference between data and model. This cost function describes a statistical distribution of errors which is like a Gaussian, but with wider wings. This is a good distribution function when the data are approximately Gaussian, but have some outliers which do not fit the Gaussian model. Another statistic is the Poisson distribution, which is particularly useful for counting data with very small average counts. In this case, we seek to maximize the probability of the data given the count rates modeled for each pixel. For even modest count rates this probability is very small because there are so many combinations of count rates. So we choose to minimize the negative of the log of that probability. In this paper we show only results of using χ^2 as the cost function.

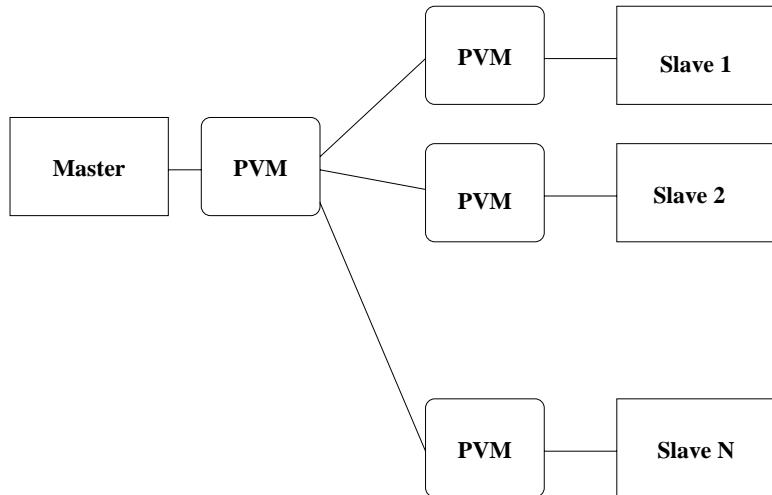


Figure 2. Process farming architecture.

3.5. The optimization function

The optimization function is the procedure for finding the minimum value of the cost function. So far we have used only the Powell minimization function. While it is a greedy algorithm, it has proved to function well enough for initial tests. We have however also demonstrated that it does occasionally fall into a local, non-optimal minimum. This demonstration was done by attempting Powell minimization with a large number of starting points. Often, different starting points would result in different minima. In this paper we chose to ignore that source of error, but acknowledge that it needs to be addressed, or at least quantified, in the future.

3.6. Computer implementation

Because of the expected large computational burden in performing the fits, we decided to parallelize the software within the so-called process-farming framework (Figure 2). In this framework a master creates computing tasks which are executed in an asynchronous fashion by processing slaves. A job transmitted to the slave consists of a list of frames to be fitted, the description of the functional form of the parameters as a function of time, and selection for cost-function and optimization function. The slave performs the computation and returns the best parameters found. The master writes these parameters to a file and sends a new task to the slave. When all tasks have been submitted once, the master begins to re-submit tasks that have not yet returned, in order to guard against slow or down machines. We have found this approach to be very efficient for the computational complexity we have exposed the system to. We have typically run the system on a network of up to 120 workstations, with up to approximately 200 processors. A computation on a 15000 frame scan can take from a few minutes to 10 hours of wall-clock time, corresponding to a total CPU usage ranging from 3 hours to 2000 hours. We also found that the network burden and computational burden on the master are very small such that we could in principle run the program on a much larger cluster (perhaps 10 or more times as large) if we had one available. The reason for the very large range of compute times is that the optimization time depends on the formulation of the problem, such as the number of parameters to fit for, the number of frames to fit, and the choice of cost- and optimization functions.

4. DISCUSSION

4.1. Reduction in the number of outliers

For both the fitting and the group delay methods there are secondary extrema in the cost function. For the fitting approach, which does not fully sample the search space it is possible to sometimes hit one of those secondary

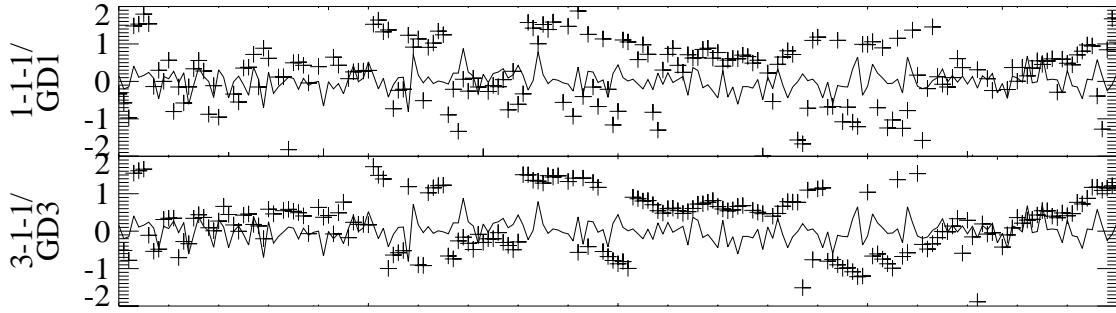


Figure 3. Fringe tracking error for various fitting schemes compared. The symbols in front of the slash on the y-axis labels are the plus symbols on the plot, and the symbols after the slash are the solid curves.

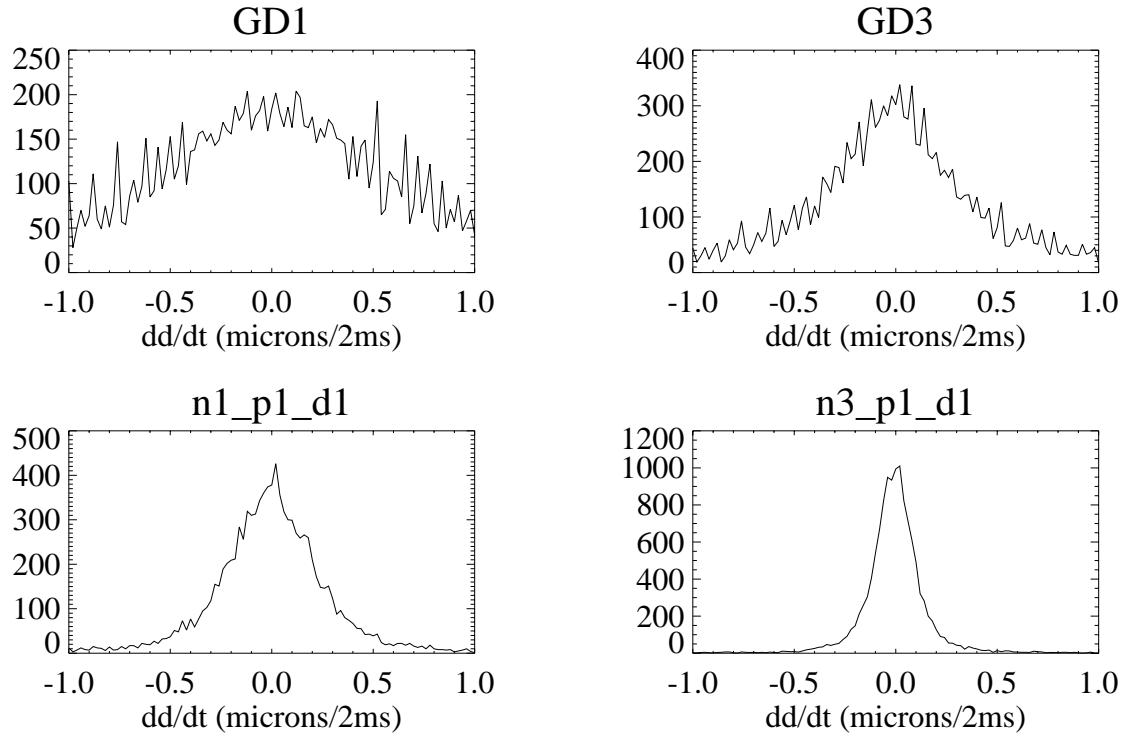


Figure 4. Histogram of the frame-to-frame difference in the fringe tracking error, derived from both the fits and the group delay method.

extrema instead of the global extremum. Furthermore, for both approach, when the count rate is very low, the true parameters - which we seek - may not be associated with the optimal extremum in the cost function due to counting statistics. When we increase the number of frames fitted simultaneously we reduce the number of outliers. Figure 3 shows the delay as a function of time for the first two-hundred frames at one fringe frequency for various fitting models. Figure 4 shows histograms of the frame-to-frame difference for an entire scan (15000 frames).

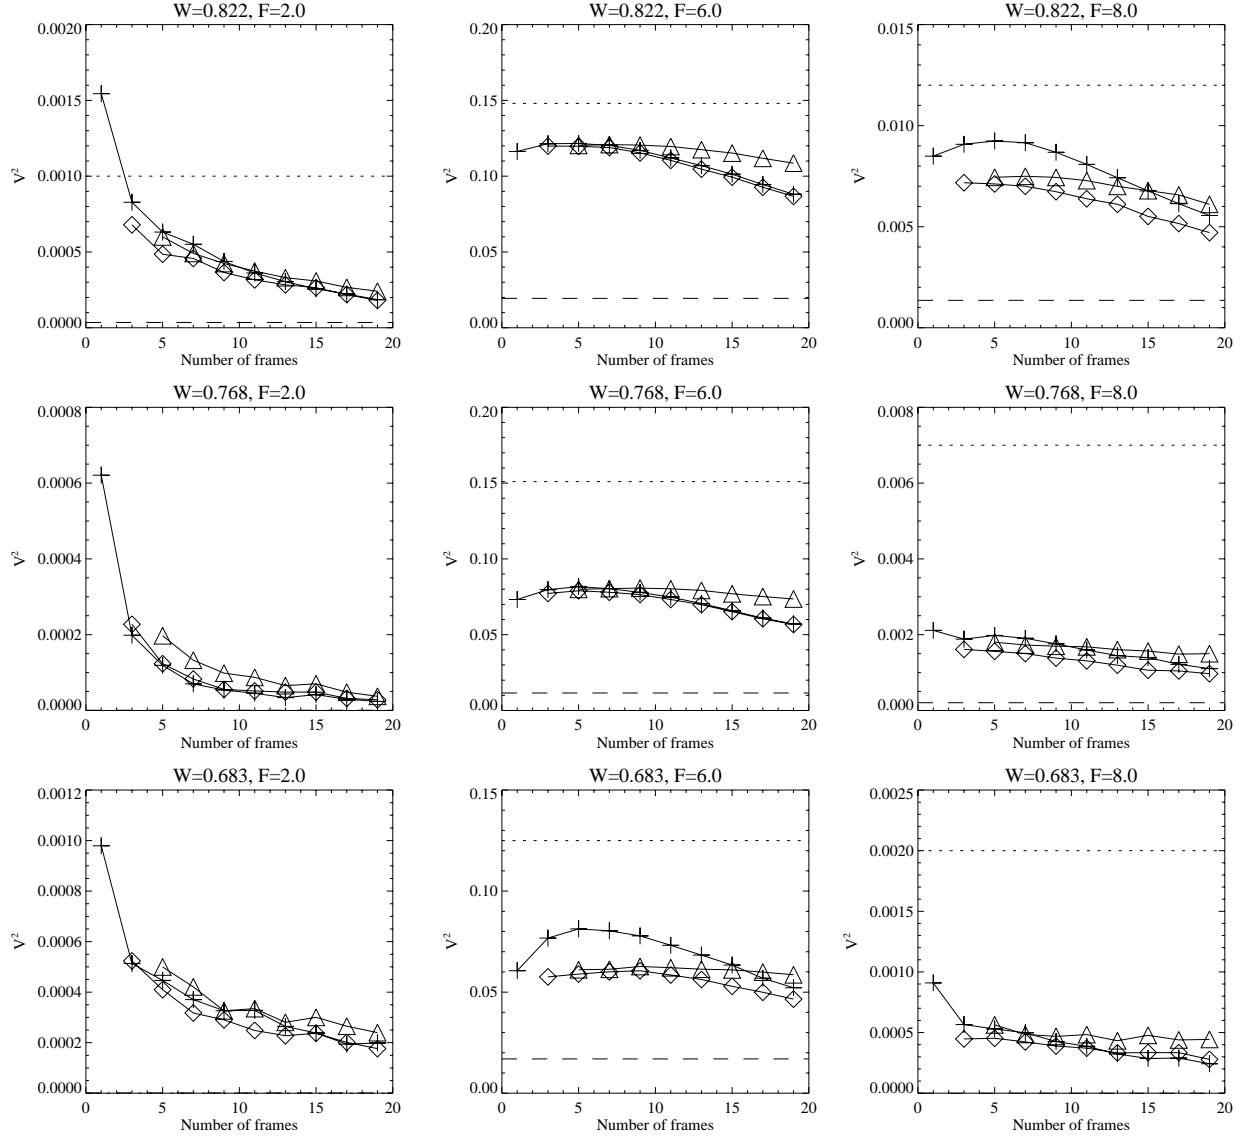


Figure 5. Comparison of different fitting models for the bright program star. Each panel shows a different wavelength/baseline combination, out of a total of 48. The vertical axis is V^2 , and the horizontal axis is the number of frames in a fit. The dotted line is V^2 obtained from a fit to the fringe power spectrum. The dashed curve is the coherently integrated V^2 using the Fourier transform group delay method with a single frame to obtain the phases. All curves are for a zeroth order polynomial (constant) for phase, ϕ . The solid curves with '+'-symbols, diamonds, and triangles, are for a zeroth order, first order, and second order polynomials for fringe tracking error, d .

4.2. Comparison of models

The fitting approach described above allows for a wide variety of expressions. Even if we restrict ourselves to one cost function (as mentioned earlier, we show only results using a χ^2 cost function), and to polynomials as a function of time for phase and delay, there are still a large number of possible expressions. We can choose the number of frames to fit simultaneously, we can choose the order of the phase polynomials, and we can choose the order of the delay polynomials. We label these as X_Y_Z , where X is the number of frames to fit simultaneously, Y is the number of phase terms, and Z is the number of delay terms. For the standard group delay calculation we use the symbol GDX , where X is the number of consecutive frames added before the calculation is performed.

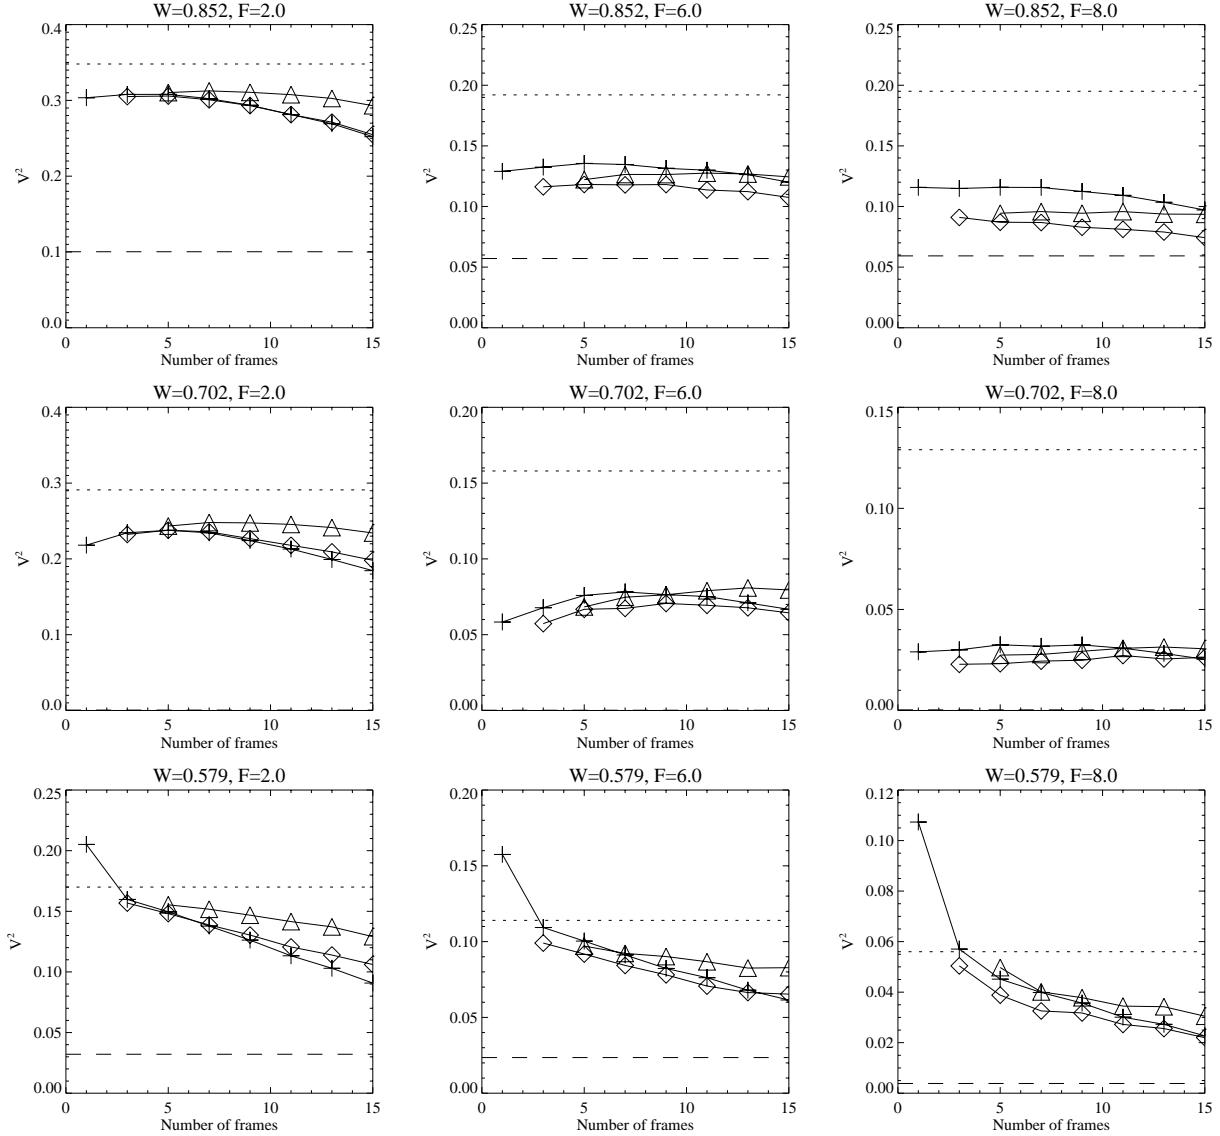


Figure 6. Comparison of different fitting models for the faint calibration star. Refer to the caption for Figure 5 for explanation.

Figure 5 shows the results for the bright star. Firstly, it can be seen that the fitting procedure does a better job at determining the phase and fringe tracking error than the group delay method, even in the case of just one frame. Furthermore, it can be seen that there is a peak in V^2 for the constant form of the fringe tracking error ('+'-symbols) near 3-7 frames, consistent with multiple frames providing improved determination. For higher-order forms of the fringe tracking error the gain is not clear, and more work is needed to understand the reason for this - one possibility is that the fringe function becomes sufficiently complicated that a non-optimal minimum is frequently found.

Figure 6 shows the results for the faint calibration star. The curves show in some cases a gain by using multiple frames, and by using higher-order forms of the fringe tracking error. While the performance is generally significantly better than the Fourier-transform group-delay method, more work is needed to understand the details of the improvements, and why fitting multiple frames does not always provide the improvement expected.

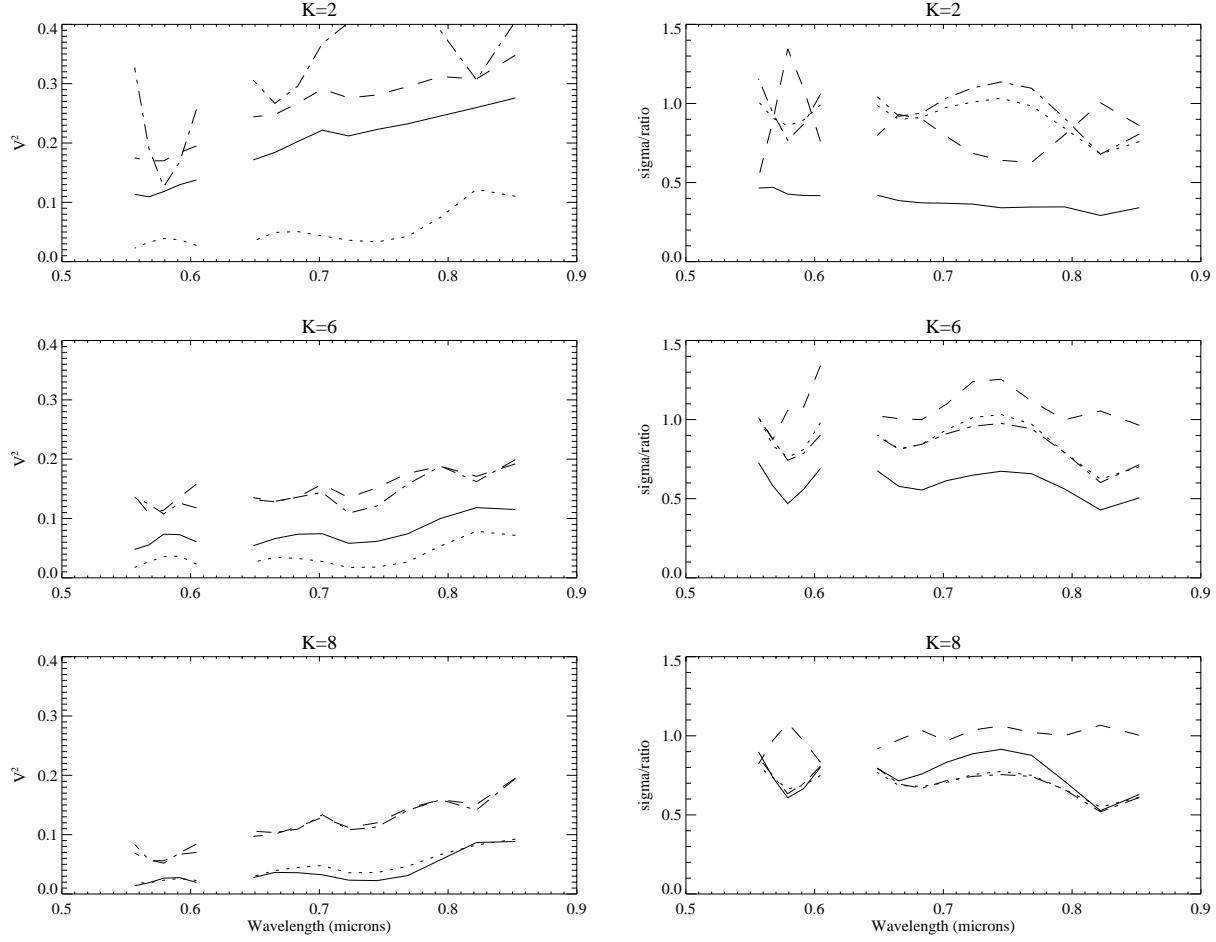


Figure 7. Illustration of the calibration process for three different baselines on the same spectrograph. Each row is for one baseline. The left-hand column shows V^2 , whereas the right-hand column shows σ , and ratios of incoherent and coherent V^2 . In the left-hand column, the dashed curve is the incoherent V^2 , the solid curve is the coherently integrated V^2 , the dotted curve is the coherently integrated V^2 bootstrapped from the other two baselines, and the dash-dotted curve is the calibrated, coherently integrated V^2 after applying Equation 8 and Equation 10. In the right-hand columns, the solid curve shows σ , the dash-dotted curve shows the bootstrapped σ , the dotted line σ between the bootstrapped V^2 and the incoherent V^2 , and the dashed curve the ratio of the incoherent V^2 and the bootstrapped V^2 .

4.3. Calibration

If the correction phase in Equation 3 is not accurately determined the coherently integrated visibility will be lower than the true instrumental visibility. If we assume a Gaussian distribution of the phase error this can be written as

$$V = V_0 e^{-\sigma^2}, \quad (8)$$

with V being the coherently integrated visibility, and V_0 being the true instrumental visibility. The problem now becomes one of determining σ . Quirrenbach et al. (1994)⁴ goes into the description of the various sources of uncertainty in some detail. A practical and precise determination of σ from first principles is, however, difficult. However, if we have two high-visibility baselines (1 and 2) and one low-visibility baseline (3) forming a triangle, we can bootstrap phase information on the low-visibility baseline from the high-visibility baseline [e.g. Hummel et al. (2003)⁵]. If at the same time we use the fringe spectrum method to incoherently determine V_0^2 on the high-visibility baselines, we can estimate σ as follows

$$\sigma_{1|2}(\lambda)^2 = \frac{\ln(V_0^2(\lambda)) - \ln(V^2(\lambda))}{2}. \quad (9)$$

Since the phase for which σ is computed comes directly from $d_k(i)$ and $\phi_k(i)$ and these quantities for the third baseline is a sum or difference of the quantities for the other two baselines, the σ^2 should also add, assuming Gaussian noise,

$$\sigma_3^2 = \sigma_1^2 + \sigma_2^2, \quad (10)$$

and σ_3 can readily be used with Equation 8 above to determine $V_{3,0}$, and calibration with a calibrator star can proceed in the usual way. It should be noted that when we are working with complex visibilities, including phases, and working in the resolved regime, the calibration requirements on V (or V^2) become much less stringent as the models are usually much more different in the resolved regime than they are in the unresolved regime.

Figure 7 shows a comparison between V^2 obtained through the incoherent method, and V^2 obtained through a coherent integration in combination with σ_1 and σ_2 obtained for the other two baselines by comparing V_{coh}^2 and V_{inc}^2 . In the left-hand column the dashed curve and dot-dashed curve must agree for the calibration to be successful. We can see that for $K = 2$ they do not agree very well, whereas they agree well for $K = 6$, and even better for $K = 8$. In the right-hand column, the dotted and dot-dashed curves should agree. The bootstrapped σ should be identical to the sigma between the bootstrapped V^2 and the incoherent V^2 . While we have some indication that the methods work, at least in two out of three cases, there is still work to be done to improve and validate the method.

4.4. Baseline phases

When coherently integrating, we obtain enough information to determine phases on individual baselines. The phase is obtained with respect to zero fringe tracking error. So the phase is the phase of the fringe at zero total delay. The atmosphere introduces an unknown phase variation which varies as a function of wavelength approximately as a quadratic. The glass in NPOI introduces an unknown phase variation which varies as a function of wavelength approximately as a cubic. Thus, phase variations up to a cubic can be due either to the astronomical source or to the instrument or the atmosphere. However, any higher-order phase variations are due to just the source. Figure 8 shows the phase and real component of the complex visibility for the bright calibration star. The zero in panel (f) corresponds well to the 180° phase change in panel (e).

4.5. Precise diameter measurements

When there is a null in the visibility function it is possible to measure the diameter of the star very precisely. Unlike measuring diameters by accurately calibrating V^2 and fitting a model to it, using the null requires very little calibration. Meisner (2003)¹⁰ demonstrated this. With NPOI we can currently measure the visibility as a function of 16 wavelengths. If there is a null in that wavelength range, we can use it, in combination with the baseline length, the angular diameter of the star at the wavelength of the null. As the Earth rotates and the projected baseline length changes, the null may occur at a different wavelength. In this way the star's diameter can be measured as a function of wavelength. If a model is fitted to it, the star's radius can be determined very accurately, including accurately enough to measure the change in radius with pulsation phase. The precision of the determination depends not on the precision of the calibration of visibilities, but rather on the precision of calibration of the wavelength scale and the number of spectral channels. Figure 9 illustrates how the diameter of a star is measured. Panels (a) and (b) show the coherently integrated X and Y for each of the 30 s scans. Panel (c) shows the phase values for all scans at all wavelength, clearly showing a 180 degree phase shift in the range 0.65 μm to 0.75 μm . Panel (d) shows the location of the null as a function of observation time ('+') and the baseline length in the same time interval. Panel (e) shows the stellar diameter as a function of wavelength, and the solid line is a model fit. As an illustration, the model fit determines the diameter to be 3.25 mas with an uncertainty of 0.0096 mas at a wavelength of 0.7 μm , or a diameter determination to better than one part in 300 with only the wavelength scale calibrated. At this level of precision we are approaching the calibration precision of the NPOI spectrographs (although there are plans and hardware in place to improve that calibration in the future).

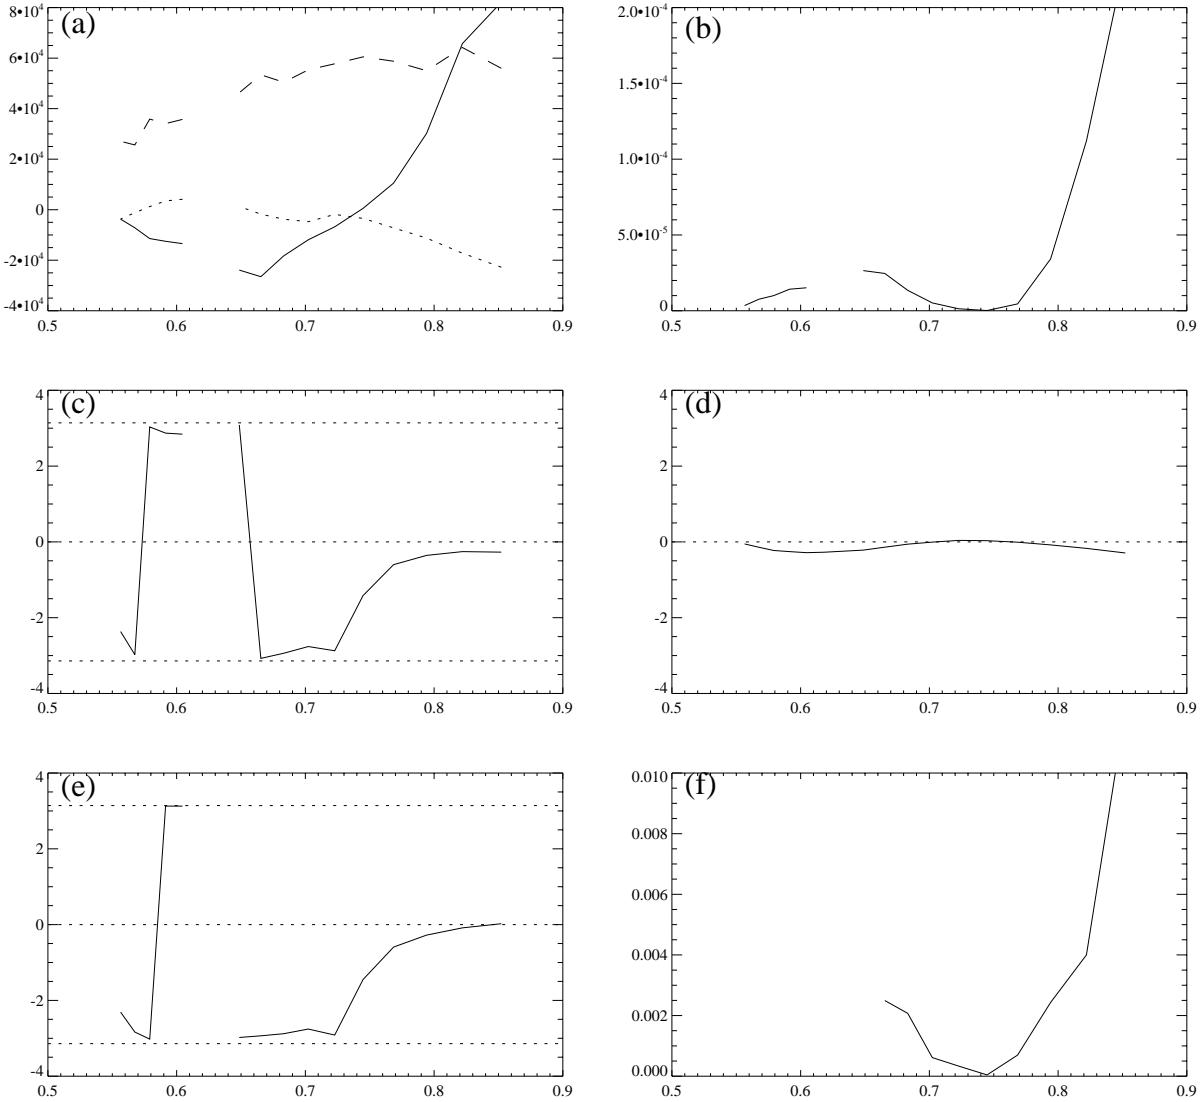


Figure 8. Illustration of the computation of baseline phases and calibration of visibilities. This is for the resolved baseline on the bright program star. (a) X (solid), Y (dotted), and N/100 (dashed), (b) uncalibrated coherently integrated V^2 obtained from X, Y, and N, (c) uncalibrated visibility phase (obtained as $\tan^{-1}(Y, X)$), (d) instrumental phase of the faint calibrator star, (e) calibrated phase for program star obtained by subtracting the phase in (d) from the phase in (c), (f) calibrated instrumental V^2 obtained through the method outlined earlier in this paper.

5. CONCLUSION

We have demonstrated a multiple-frame fringe fitting method that can be used to determine the fringe phase as a function of time and wavelength, even for situations where the count rate (S/N) is sufficiently small that a large number of outliers are produced (i.e. when the wrong peak in GD power spectrum is detected). When combined with a bootstrapping approach we can obtain coherently integrated visibilities even for baselines with very faint fringes. We have illustrated a method for calibrating coherently averaged visibility amplitudes, although more work is needed to better understand the process and uncertainties. We have shown that we can measure single baseline visibility phases accurately, and that we can use measurements of visibility nulls to make very precise determinations of stellar diameters.

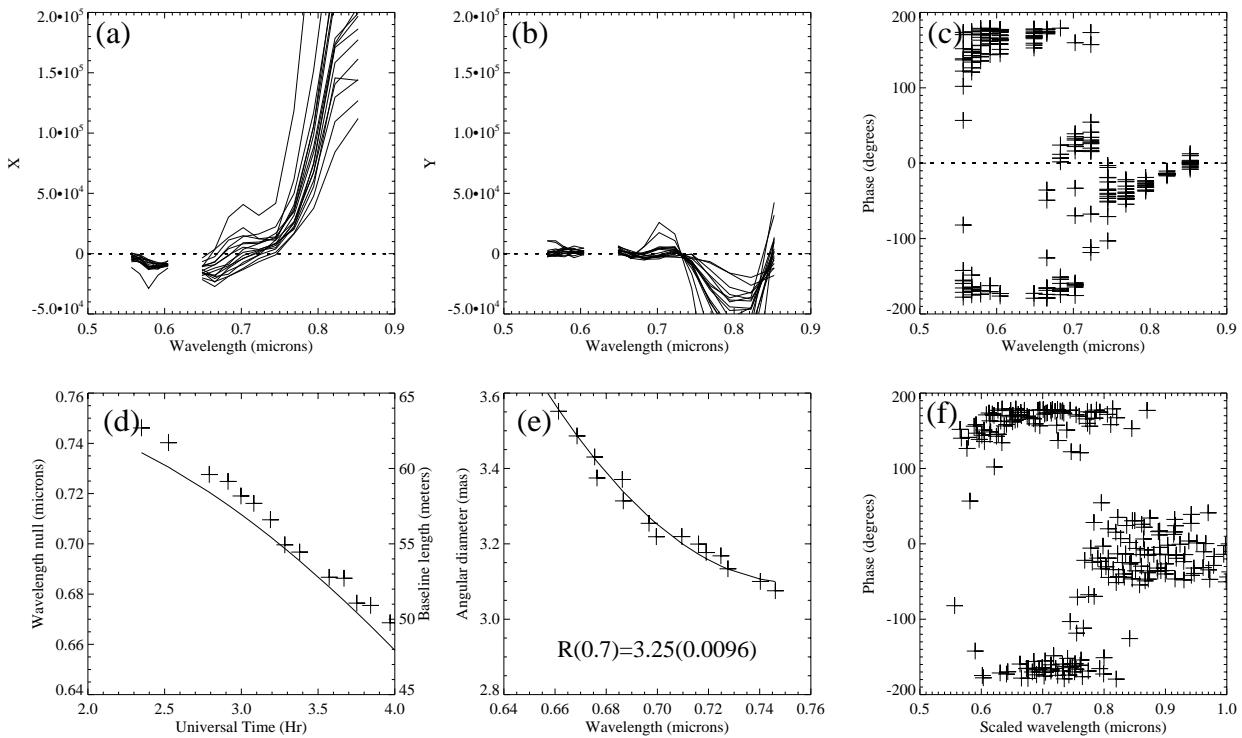


Figure 9. Illustration of the procedure for computing very precise stellar diameters. The data set consists of many individual 30 second scans recorded through a night. (a) X for each scan, (b) Y for each scan, (c) uncalibrated phase for each scan, (d) wavelength of each zero crossing in X for each scan, and baseline length as a function of time, (e) diameter as a function of wavelength, (f) phase corrected for baseline length.

REFERENCES

1. M. M. Colavita, "Fringe visibility estimators for the palomar testbed interferometer," *Pub. Astron. Soc. Pac.* **111**, pp. 111–117, 1999.
2. J. A. Meisner, "Coherent integration of fringe visibility employing probabilistic determinations of atmospheric delay," *Proc. SPIE* **3350**, pp. 294–308, 1998.
3. J. A. Meisner, "Coherent integration of fringe visibility: a generalized approach," *Proc. SPIE* **4006**, pp. 1068–1082, 2000.
4. A. Quirrenbach, D. Mozurkewich, D. F. Buscher, C. A. Hummel, and J. T. Armstrong, "Phase-referenced visibility averaging in optical long-baseline interferometry," *Astron. Astrophys.* **286**, pp. 1019–1027, 1994.
5. C. A. Hummel, D. Mozurkewich, J. A. Benson, and M. Wittkowski, "Coherent integration using phase bootstrapping," *Proc. SPIE* **4838**, pp. 1107–1114, 2003.
6. B. F. Lane and M. M. Colavita, "Phase-referenced stellar interferometry at the palomar testbed interferometer," *Astron. J.* **125**, pp. 1623–1628, 2003.
7. M. Wittkowski and C. A. Hummel, "Interferometric measurements of stellar intensity profiles," *Proc. SPIE* **4838**, pp. 210–220, 2003.
8. J. T. Armstrong, D. Mozurkewich, L. J. Rickard, D. J. Hutter, J. A. Benson, P. F. Bowers, N. M. Elias II, C. A. Hummel, K. J. Johnstron, D. F. Buscher, J. H. Clark III, L. Ha, L.-C. Ling, N. M. White, and R. S. Simon, "The navy prototype optical interferometer," *Astrophys. J.* **496**, pp. 550–571, 1998.
9. M. Born and E. Wolf, *Principles of Optics*, 6 edn., Pergamon Press, Oxford, UK, 1980.
10. J. Meisner, "Scientific and technical results from VINCI using coherent estimation of fringe visibility," *Astrophysics and Space Science* **286**, pp. 119–127, 2003.